arXiv:1405.0329v1 [cs.DM] 2 May 2014

Forbidden Induced Subgraphs of Normal Helly Circular-Arc Graphs: Characterization and Detection*

Yixin Cao† Luciano N. Grippo‡ Martín D. Safe‡

May 5, 2014

Abstract

A normal Helly circular-arc graph is the intersection graph of arcs on a circle of which no three or less arcs cover the whole circle. Lin, Soulignac, and Szwarcfiter [Discrete Appl. Math. 2013] characterized circular-arc graphs that are not normal Helly circular-arc graphs, and used it to develop the first recognition algorithm for this graph class. As open problems, they ask for the forbidden induced subgraph characterization and a direct recognition algorithm for normal Helly circular-arc graphs, both of which are resolved by the current paper. Moreover, when the input is not a normal Helly circular-arc graph, our recognition algorithm finds in linear time a minimal forbidden induced subgraph as certificate.

Keywords: certifying algorithms, holes, interval models, (minimal) forbidden induced subgraphs, (normal, Helly) circular-arc models.

1 Introduction

This paper will be only concerned with undirected and simple graphs. A graph is a *circular-arc graph* if its vertices can be assigned to arcs on a circle such that two vertices are adjacent if and only if their corresponding arcs intersect. Such a set of arcs is called a *circular-arc model* of this graph. If there is some point on the circle that is not covered by any arc in the model, then the graph is an *interval graph*, and it can also be represented by a set of intervals on the real line, which is called an *interval model*. Circular-arc graphs and interval graphs are two of the most famous intersection graph classes, and both have been studied intensively for decades. However, in contrast to interval graphs, our understanding of circular-arc graphs is far limited, and to date some fundamental problem remains unsolved.

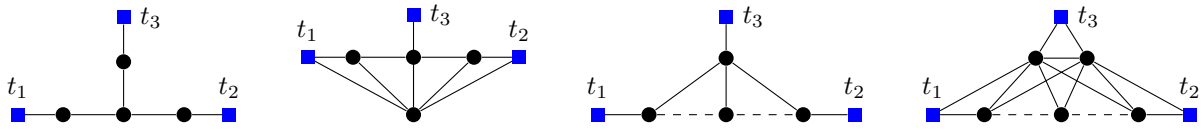

(a) long claw (b) whipping top (c) † (d) ‡

Figure 1: Chordal minimal forbidden induced graphs.

One fundamental combinatorial problem on a graph class is its characterization by forbidden induced subgraphs. For example, Lekkerkerker and Boland [8] showed in 1962 that a graph is an interval graph if and only if it contains neither a hole (i.e., a induced cycle of length at least four) nor any graph in Fig. 1 as an induced

*Preliminary results of this paper appeared in the proceedings of SBPO 2012 [4] and FAW 2014 [1].

†Institute for Computer Science and Control, Hungarian Academy of Sciences. Email: yixin@sztaki.hu. Supported by the European Research Council (ERC) under the grant 280152 and the Hungarian Scientific Research Fund (OTKA) under the grant NK105645.

‡Instituto de Ciencias, Universidad Nacional de General Sarmiento, Los Polvorines, Buenos Aires, Argentina. Email:lgrippo@ungs.edu.ar, msafe@ungs.edu.ar. Partially supported by CONICET PIP 11220120100450CO and ANPCyT PICT 2012-1324 grants.

subgraph. Recall that holes are the forbidden induced subgraphs of *chordal graphs*, which are the intersection of subtrees of a tree. In contrast, since it was first asked by Hadwiger et al. [5] in 1964, all efforts attempting to characterize circular-arc graphs by forbidden induced subgraphs have succeeded only partially. Tucker made the most significant contribution to the study of the class of circular-arc graphs and its subclasses, which includes the forbidden induced subgraph characterizations of both unit circular-arc graphs (i.e., a graph with a circular-arc model where every arc has the same length) and proper circular-arc graphs (i.e., a graph with a circular-arc model where no arc properly contains another) [17]; we will see more later. There is a similar line of research for other proper subclasses of circular-arc graphs, which aims at determining their forbidden induced subgraphs or some other kinds of obstructions; for this we refer to the surveys of Lin and Soulignac [11] and Durán et al. [3] and references therein.

One fundamental algorithmic problem on a graph class is its recognition, i.e., to efficiently decide whether a given graph belongs to this class or not. For intersection graph classes, all recognition algorithms known to the authors provide an intersection model when the membership is asserted. Most of them, on the other hand, simply return "NO" for non-membership, while one might also want some verifiable *certificate* for some reason [14]. A recognition algorithm is *certifying* if it provides both positive and negative certificates. There are different forms of negative certificates, while a minimal forbidden induced subgraph is arguably the simplest and most preferable of them [6]. Kratsch et al. [7] reported a certifying recognition algorithm for interval graphs, which in linear time returns either an interval model of an interval graph or a forbidden induced subgraph for a non-interval graph. Although its returned forbidden induced subgraph is not necessarily minimal, a minimal one can be easily retrieved from it (see also [12] for another approach). Likewise, a hole can be detected from a non-chordal graph in linear time [16]. On the other hand, although a circular-arc model of a circular-arc graph can be produced in linear time [13], it remains a challenging open problem to find a negative certificate for a non-circular-arc graph.

The complication of circular-arc graphs may be attributed to two special intersection patterns of circular-arc models that are not possible in interval models. The first is two arcs intersecting in both ends, and a circular-arc model is called *normal* if no such pair exists. The second is a set of arcs intersecting pairwise but containing no common point, and a circular-arc model is called *Helly* if no such set exists. Normal and Helly circular-arc models are precisely those without three or less arcs covering the whole circle [15, 10]. A graph that admits such a model is called a *normal Helly circular-arc graph*. In particular, all interval graphs are normal Helly circular-arc graphs.

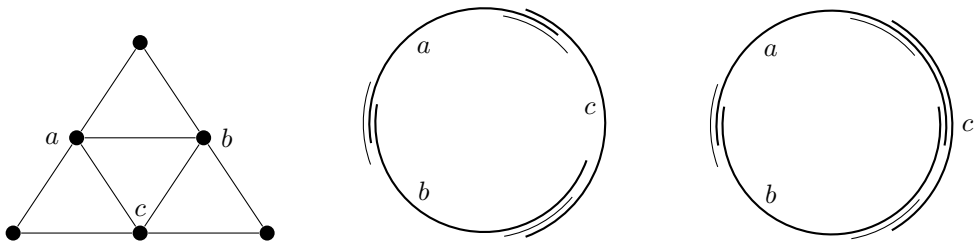

(a) A circular-arc graph G (b) A normal model of G (c) A Helly model of G

Figure 2: a non-normal Helly circular-arc graph and its circular-arc models

A word of caution is worth on the definition of normal Helly circular-arc graphs. One graph might admit both a normal circular-arc model and a Helly circular-arc model but not a normal and Helly circular-arc model. For example, see the graph and its models in Fig. 2. The fact that the model of Fig. 2b (resp., Fig. 2c) is not Helly (resp., normal) can be evidenced by arc set $\{a, b, c\}$ (resp., $\{a, b\}$). One may want to verify that arranging a normal and Helly circular-arc model for this graph is out of the question. This example convinces us that the set of normal Helly circular-arc graphs is *not* equivalent to the intersection of normal circular-arc graphs and Helly circular-arc graphs, but a proper subset of it.

Let us mention some previous work related to normal Helly circular-arc graphs. Tucker [18] gave an algorithm that outputs a proper coloring of any given normal Helly circular-arc graph using at most $3\omega/2$ colors, where ω denotes the size of a maximum clique. Note that by the Helly property, ω is equivalent to the maximum number of arcs covering a single point on the circle. This is tight as any odd hole, which has $\omega = 2$ and needs at least

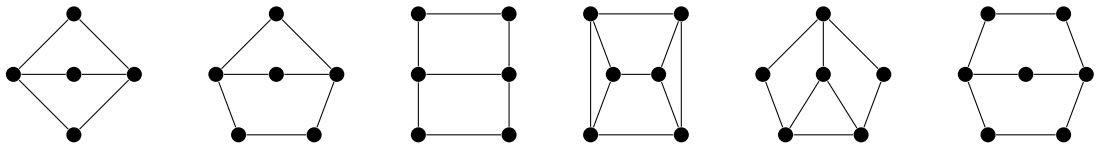

(a) $K_{2,3}$ (b) twin-C_5 (c) domino (d) $\overline{C_6}$ (e) FIS-1 (f) FIS-2

Figure 3: Non-chordal and finite minimal forbidden induced graphs.

three colors, is a normal Helly circular-arc graph. In the study of convergence of circular-arc graphs under the clique operator, Lin et al. [9] observed that normal Helly circular-arc graphs arose naturally. They then [10] undertook a systematic study of normal Helly circular-arc graphs as well as its subclass. Their results include a partial characterization of normal Helly circular-arc graphs by forbidden induced subgraph (more specifically, those restricted to Helly circular-arc graphs), and a linear-time recognition algorithm (by calling a recognition algorithm for circular-arc graphs). As open problems, they ask for determining the remaining minimal forbidden induced subgraphs, and designing a direct recognition algorithm, both of which are resolved by the current paper.

The first main result of this paper is a complete characterization of normal Helly circular-arc graphs by forbidden induced subgraphs. A wheel (resp., C^*) comprises a hole and another vertex completely adjacent (resp., nonadjacent) to it.

Theorem 1.1. *A graph is a normal Helly circular-arc graph if and only if it contains no C^*, wheel, or any graph depicted in Figs. 1 and 3.*

It is easy to use the definition to verify that a normal Helly circular-arc graph is chordal if and only if it is an interval graph. An interval model is always a normal and Helly circular-arc model, but an interval graph might have circular-arc model that is neither normal nor Helly, e.g., consider K_4. For non-chordal graphs we have:

Proposition 1.2 ([15, 10])**.** *If a normal Helly circular-arc graph G is not chordal, then every circular-arc model of G is normal and Helly.*

These observations inspire us to recognize normal Helly circular-arc graphs as follows. If the input graph is chordal, it suffices to check whether it is an interval graph. Otherwise, we try to build a circular-arc model of it, and if we succeed, verify whether the model is normal and Helly. Lin et al. [10] showed that this approach can be implemented in linear time. Moreover, if there exists a set of at most three arcs covering the circle, then their algorithm returns it as a certificate. This algorithm, albeit conceptually simple, suffers from twofold weakness. First, it needs to call some recognition algorithm for circular-arc graphs, while all known algorithms are extremely complicated. Second, it is very unlikely to deliver a negative certificate in general.

The second main result of this paper is the following direct certifying algorithm for recognizing normal Helly circular-arc graphs, which would be desirable for both efficiency and the detection of negative certificates. From now on, unless otherwise stated, whenever we refer to a "minimal forbidden induced subgraph" it should be understood a minimal forbidden induced subgraph for the class of normal Helly circular-arc graphs. We use $n := |V(G)|$ and $m := |E(G)|$ throughout.

Theorem 1.3. *There is an $O(n + m)$-time algorithm that given a graph G, either constructs a normal and Helly circular-arc model of G, or finds a minimal forbidden induced subgraph of G.*

It is clear that each graph specified in Theorem 1.1 is a minimal forbidden induced subgraph. First, every graph in Fig. 1 is chordal but non-interval graph, and thus cannot be a normal Helly circular-arc graph. Second, a C^* is not a circular-arc graph, while a wheel cannot be arranged without three or less arcs covering the circle. Third, every graph in Fig. 3 has only a small number of vertices and can be easily checked. Therefore, to prove Theorem 1.1, it suffices to show that a graph containing none of them is a normal Helly circular-arc graph. That fact was actually proved in [4], but the resulting proof of Theorem 1.1 given there does not provide a linear-time procedure to find the corresponding forbidden induced subgraphs when the graph is not a normal Helly circular-arc graph. Since the algorithm we use to prove Theorem 1.3 always finds such a subgraph in this case, Theorem 1.1 follows from the correctness proof of our algorithm as a corollary.

Let us briefly discuss the basic idea behind the way we deal with a non-chordal graph G. If G is a normal Helly circular-arc graph, then for any vertex v of G, both $N[v]$ and its complement induce nonempty interval subgraphs. The main technical difficulty is how to combine interval models for them to make a circular-arc model of G. For this purpose we build an auxiliary graph $\mho(G)$ by taking two identical copies of $N[v]$ and appending them to the two ends of $G - N[v]$ respectively. The shape of symbol $\mho$ is a good hint for understanding the structure of the auxiliary graph. We show that $\mho(G)$ is an interval graph and more importantly, a circular-arc model of G can be produced from an interval model of $\mho(G)$. On the other hand, if G is not a normal Helly circular-arc graph, then $\mho(G)$ cannot be an interval graph. In this case we use the following procedure to obtain a minimal forbidden induced subgraph of G.

Theorem 1.4. *Given a minimal non-interval induced subgraph of $\mho(G)$, we can in $O(n + m)$ time find a minimal forbidden induced subgraph of G.*

The crucial idea behind our certifying algorithm is a novel correlation between normal Helly circular-arc graphs and interval graphs, which can be efficiently used for algorithmic purpose. This was originally proposed in the detection of small forbidden induced subgraph of interval graphs [2], i.e., the opposite direction of the current paper. In particular, in [2] we have used a similar definition of the auxiliary graph and pertinent observations. However, the main structures and the procedures for the detection of forbidden induced subgraphs divert completely. For example, the most common forbidden induced subgraphs in [2] are 4- and 5-holes, which, however, are allowed in normal Helly circular-arc graphs. This means that the interaction between $N[v]$ and $G - N[v]$ are far more subtle, and thus the detection of minimal forbidden induced subgraphs in the current paper is significantly more complicated than that of [2].

2 The recognition algorithm

All graphs are stored as adjacency lists. We use the customary notation $v \in G$ to mean $v \in V(G)$, and $u \sim v$ to mean $uv \in E(G)$. The *degree* of a vertex v is defined by $d(v) := |N(v)|$, where $N(v)$, called the *neighborhood* of v, comprises all vertices u such that $u \sim v$. The *closed neighborhood* of v is defined by $N[v] := N(v) \cup \{v\}$. For a vertex set U, its closed neighborhood and neighborhood are defined by $N[U] := \bigcup_{v \in U} N[v]$ and $N(U) := N[U] \backslash U$, respectively. Exclusively concerned with induced subgraphs, we use F to denote both a subgraph and its vertex set.

Consider a circular-arc model $\mathcal{A}$. If every point of the circle is contained in some arc in $\mathcal{A}$, then we can find an inclusion-wise minimal set X of arcs that cover the entire circle. If $\mathcal{A}$ is normal and Helly, then X consists of at least four vertices and thus corresponds to a hole. Therefore, a normal Helly circular-arc graph G is chordal if and only if it is an interval graph, for which it suffices to call the algorithms of [7, 12]. We are hence focused on graphs that are not chordal. We call the algorithm of Tarjan and Yannakakis [16] to detect a hole H.

Proposition 2.1. *Let H be a hole of a circular-arc graph G. In any circular-arc model of G, the union of arcs for H covers the whole circle. In other words, $N[H] = V(G)$.*

The indices of vertices in the hole $H = (h_0 h_1 \dots h_{|H|-1} h_0)$ should be understood as modulo $|H|$, e.g., $h_{-1} = h_{|H|-1}$. By Proposition 2.1, every vertex should have neighbors in H. We use $N_H[v]$ as a shorthand for $N[v] \cap H$, regardless of whether $v \in H$ or not. We start from characterizing $N_H[v]$ for every vertex v: we specify some forbidden structures not allowed to appear in a normal Helly circular-arc graph, and more importantly, we show how to find a minimal forbidden induced subgraph if one of these structures exists. The fact that they are forbidden can be easily seen from the definition of normal and Helly and Proposition 2.1, and hence the proofs given below will focus on the detection of minimal forbidden induced subgraphs.

Lemma 2.2. *For every vertex v, we can in $O(d(v))$ time find either a proper sub-path of H induced by $N_H[v]$, or a minimal forbidden induced subgraph.*

Proof. We pre-allocate a list IND of $d(v)$ slots, initially all empty. For each neighbor of v, if it is h_i, then add i into the next empty slot of IND. After all neighbors of v have been checked, we shorten IND by removing empty slots from the end, which leaves $|N_H[v]|$ slots. If $|N_H[v]|$ is 0 or $|H|$, then we return H and v as a C^* or wheel. In the remaining case, $N_H[v]$ is a nonempty and proper subset of H. We radix sort IND; let p and q be its first and last elements respectively.

Starting from the first element, we traverse IND to the end for the first i such that $\mathsf{IND}[i+1] > \mathsf{IND}[i]+1$. If no such i exists, then we return $(h_p \cdots h_q)$ as the path. In the remaining cases, we may assume that we have found the i; let $p_1 := \mathsf{IND}[i]$ and $p_2 := \mathsf{IND}[i+1]$. We continue to traverse from $i+1$ to the end of IND for the first j such that $\mathsf{IND}[j+1] > \mathsf{IND}[j]+1$. This step has three possible outcomes: (1) if j is found, then $p_3 := \mathsf{IND}[j]$ and $p_4 := \mathsf{IND}[j+1]$; (2) if no such j is found, and at least one of $q < |H|-1$ and $p > 0$ holds, then $p_3 := q$ and $p_4 := p + |H|$; and (3) otherwise ($p = 0$, $q = |H|-1$, and j is not found). In the third case, we return $(h_{p_2} \cdots h_{|H|-1} h_0 \cdots h_{p_1})$ as the path induced by $N_H[v]$. In the first two cases, p_3 and p_4 are defined, and $p_4 > p_3 + 1$. In other words, we have two nontrivial sub-paths, $(h_{p_1} h_{p_1+1} \dots h_{p_2})$ and $(h_{p_3} h_{p_3+1} \dots h_{p_4})$, of H such that v is adjacent to their ends but none of their inner vertices.

If $p_2 - p_1 > 3$, then we return $(v h_{p_3} h_{p_3+1} \dots h_{p_4} v)$ and h_{p_1+2} as a C^*. Likewise, if $v \not\sim h_\ell$ for some ℓ with $p_2 + 1 < \ell < p_1 - 1 + |H|$, then we return $(v h_{p_1} h_{p_1+1} \dots h_{p_2} v)$ and h_ℓ as a C^*; note this must hold true when v is adjacent to both h_{p_1-1} and h_{p_2+1}. Hence we may assume $2 \leqslant p_2 - p_1 \leqslant 3$, and without loss of generality, $v \not\sim h_{p_1-1}$.

If $p_2 - p_1 = 2$, then we return *(1)* $H \cup \{v\}$ as a $K_{2,3}$ when $|H| = 4$; *(2)* $H \cup \{v\}$ as a twin-C_5 when $|H| = 5$ and $|N_H[v]| = 2$; *(3)* $H \cup \{v\}$ as an FIS-1 when $|H| = 5$ and $|N_H[v]| = 3$; or *(4)* $\{h_{p_1-2}, h_{p_1-1} \cdots, h_{p_2}, v\}$ as a domino when $|H| > 5$. Otherwise, $p_2 - p_1 = 3$, and we return *(1)* $H \cup \{v\}$ as a twin-C_5 when $|H| = 5$; *(2)* $H \cup \{v\}$ as an FIS-2 when $|H| = 6$ and $v \not\sim h_{p_2+1}$; *(3)* $(v h_{p_1} h_{p_1-1} h_{p_1-2} v)$ and h_{p_2-1} as a C^* when $|H| = 6$ and $v \sim h_{p_2+1}$; or *(4)* $(v h_{p_1} h_{p_1-1} h_{p_1-2} v)$ and h_{p_2-1} as a C^* when $|H| > 6$.

The construction of IND takes $O(d(v))$ time. In the same time we can traverse it to find indices p_1, p_2, p_3, p_4. The rest uses constant time. This concludes the time analysis and completes the proof. □

We designate the ordering $h_0, h_1, h_2, \cdots$ of traversing H as *clockwise*, and the other *counterclockwise*. In other words, edges h_0h_1 and h_0h_{-1} are clockwise and counterclockwise from h_0, respectively. Now let P be the path induced by $N_H[v]$. We can assign a direction to P in accordance to the direction of H, and then we have clockwise and counterclockwise ends of P. For technical reasons, we assign canonical indices to the ends of the path P as follows.

Definition 1. *For each vertex $v \in G$, we denote by $\mathtt{first}(v)$ and $\mathtt{last}(v)$ the indices of the counterclockwise and clockwise, respectively, ends of the path induced by $N_H[v]$ in H satisfying*

- $-|H| < \mathtt{first}(v) \leqslant 0 \leqslant \mathtt{last}(v) < |H|$ *if* $h_0 \in N_H[v]$*; or*
- $0 < \mathtt{first}(v) \leqslant \mathtt{last}(v) < |H|$*, otherwise.*

It is possible that $\mathtt{last}(v) = \mathtt{first}(v)$, when $|N_H[v]| = 1$. In general, $\mathtt{last}(v) - \mathtt{first}(v) = |N_H[v]| - 1$, and $v = h_i$ or $v \sim h_i$ for each i with $\mathtt{first}(v) \leqslant i \leqslant \mathtt{last}(v)$. The indices $\mathtt{first}(v)$ and $\mathtt{last}(v)$ can be easily retrieved from Lemma 2.2, with which we can check the adjacency between v and any vertex $h_i \in H$ in constant time. Now consider the neighbors of more than one vertices in H.

Lemma 2.3. *Given a pair of adjacent vertices u, v such that $N_H[u]$ and $N_H[v]$ are disjoint, then in $O(n+m)$ time we can find a minimal forbidden induced subgraph.*

Proof. Clearly, neither of u and v can be in H. We may assume both $N_H[u]$ and $N_H[v]$ induce proper sub-paths; otherwise we can call Lemma 2.2. They partition H into four sub-paths, two of which are induced by $N_H[u]$ and $N_H[v]$. Denote by P_1 and P_2 the other two sub-paths; their ends are adjacent to u and v respectively, while their inner vertices, if any, are adjacent to neither u nor v..

Assume first that both P_1 and P_2 are of length 1, then $|H| = |N_H[u]| + |N_H[v]|$. If u is adjacent to a single vertex h_i in H, (noting that $|N_H[v]| \geqslant 3$,) then we return $\{h_i, h_{i-1}, h_{i+1}, u, v\}$ as a $K_{2,3}$. A symmetric argument applies when $|N_H[v]| = 1$. If $|N_H[u]| = |N_H[v]| = 2$, then we return $H \cup \{u, v\}$ as a $\overline{C_6}$. It must be in some case above if $|H| = 4$, and henceforth we assume $|H| > 4$. If u is adjacent to only h_i and h_{i+1} in H, (noting that $|N_H[v]| \geqslant 3$,) then we return $\{h_{i-1}, h_i, h_{i+1}, h_{i+2}, u, v\}$ as an FIS-1. A symmetric argument applies when $|N_H[v]| = 2$. Now that both $|N_H[u]|$ and $|N_H[v]|$ are at least 3, we return $P_1 \cup P_2 \cup \{u, v\}$ as a domino.

Assume now that, without loss of generality, P_2 is nontrivial. We can return $(vuP_1v)+h_{\mathtt{last}(v)+1}$ (when both $|N_H[u]| > 1$ and $|N_H[v]| > 1$) or (vuP_1v) and $h_{\mathtt{last}(v)+2}$ (when the length of P_1 is longer than 3) as a C^*. A symmetric argument applies when $\mathtt{first}(v) - \mathtt{last}(u) > 3$. In the remaining cases, we assume without loss of generality, $|N_H[u]| = 1$, and both paths P_1 and P_2 contain at most 4 vertices. Consequently, $|H \setminus N_H[v]| \leqslant 5$.

If P_1 is also nontrivial, then $|H \setminus N_H[v]| \geqslant 3$. We return $\{u, v, h_{\texttt{first}(v)-1}, h_{\texttt{last}(v)+1}\} \cup N_H[v]$ as a † when $|N_H[v]| > 1$. Now that $|N_H[v]| = 1$, then $|H| \leqslant 6$, and we return *(1)* $(H \setminus N_H[v]) \cup \{u, v\}$ as a long claw when $|H| = 6$; *(2)* $H \cup \{u, v\}$ as a twin-C_5 when $|H| = 4$; or *(3)* $H \cup \{u, v\}$ as a FIS-2 when $|H| = 5$. In the final case, P_1 is trivial but P_2 is nontrivial, which means that neither u nor v is adjacent to $h_{\texttt{first}(u)-1}$. If $v \sim h_{\texttt{first}(u)-2}$, then we return $\{h_{\texttt{first}(u)-2}, h_{\texttt{first}(u)-1}, h_{\texttt{first}(u)}, h_{\texttt{first}(u)+1}, u, v\}$ as *(1)* an FIS-1 when $|H| = 4$; or *(2)* a twin-C_5 when $|H| > 4$. If $v \not\sim h_{\texttt{first}(u)-2}$, then we return *(1)* $H \cup \{u, v\}$ as a domino when $|H| = 4$; or *(2)* (uvh_1h_0u) and h_{-2} as a C^* when $|H| > 4$. This procedure enters only one case, which is decided only by $N_H[u]$ and $N_H[v]$. Therefore, it can be done in $O(n + m)$ time. □

Lemma 2.4. *Given a set U of two or three pairwise adjacent vertices such that*

1) $\bigcup_{u \in U} N_H[u] = H$*; and*

2) for every $u \in U$, each end of $N_H[u]$ is adjacent to at least two vertices in U,

then we can in $O(n + m)$ time find a minimal forbidden induced subgraph.

Proof. Consider first that U contains only two vertices v_1 and v_2. The $(h_{\texttt{first}(v_1)}, h_{\texttt{last}(v_1)})$-path whose inner vertices are nonadjacent to v_1 makes a hole with v_1. This hole is completely adjacent to v_2, and thus we return a wheel.

Consider then $U = \{v_1, v_2, v_3\}$. We may assume that no two vertices of U satisfy the condition of the lemma, as otherwise we are in the previous case. Without loss of generality, assume that $h_{\texttt{last}(v_1)} \in N[v_2]$, and then $h_{\texttt{first}(v_2)} \in N[v_1]$. The $(h_{\texttt{first}(v_1)}, h_{\texttt{last}(v_2)})$-path whose inner vertices are adjacent to neither v_1 or v_2 makes a hole with v_1 and v_2. By assumption, v_3 is adjacent to every vertex in the hole, and thus we return a wheel. □

Let $T := N[h_0]$ and $\overline{T} := V(G) \setminus T$. As we have alluded to earlier, we want to duplicate T and append them to different sides of $\overline{T}$. Each edge between $v \in T$ and $u \in \overline{T}$ will be carried by only one copy of T, and this is determined by its direction specified as follows. We may assume that none of the Lemmas. 2.2, 2.3, and 2.4 applies to v or/and u, as otherwise we can terminate the algorithm by returning the forbidden induced subgraph found by them. As a result, u is adjacent to either $\{h_{\texttt{first}(v)}, \cdots, h_{-1}\}$ or $\{h_1, \cdots, h_{\texttt{last}(v)}\}$ but not both. The edge uv is said to be clockwise from T if $u \sim h_i$ for $1 \leqslant i \leqslant \texttt{last}(v)$, and counterclockwise otherwise. Let E_c (resp., E_{cc}) denote the set of clockwise (resp., counterclockwise) edges from T, and let T_c (resp., T_{cc}) denote the subsets of vertices of T that are incident to edges in E_c (resp., E_{cc}). Note that $\{E_{cc}, E_c\}$ partitions edges between T and $\overline{T}$, but a vertex in T might belong to both T_{cc} and T_c, or neither of them. We have now all the details for the definition of the auxiliary graph $\mathcal{U}(G)$.

Definition 2. *The vertex set of $\mathcal{U}(G)$ consists of $\overline{T} \cup L \cup R \cup \{w\}$, where L and R are distinct copies of T, i.e., for each $v \in T$, there are a vertex v^l in L and another vertex v^r in R, and w is a new vertex distinct from $V(G)$. For each edge $uv \in E(G)$, we add to the edge set of $\mathcal{U}(G)$*

- *an edge uv if neither u nor v is in T;*
- *two edges u^lv^l and u^rv^r if both u and v are in T; or*
- *an edge uv^l or uv^r if $uv \in E_c$ or $uv \in E_{cc}$ respectively ($v \in T$ and $u \in \overline{T}$).*

Finally, we add an edge wv^l for every $v \in T_{cc}$.

Lemma 2.5. *The numbers of vertices and edges of $\mathcal{U}(G)$ are upper bounded by $2n$ and $2m$ respectively. Moreover, an adjacency list representation of $\mathcal{U}(G)$ can be constructed in $O(n + m)$ time.*

Proof. The vertices of the auxiliary graph $\mathcal{U}(G)$ include $\overline{T}$, two copies of T, and w. So the number of vertices is $2|T| + |\overline{T}| + 1 = |V(G)| + |T| + 1 \leqslant 2n$. In $\mathcal{U}(G)$, there are two edges derived from every edge of $G[T]$ and one edge from every other edge of G. All other edges are incident to w, and there are T_{cc} of them. Therefore, the number of edges is $|E(G)| + |E(G[T])| + |T_{cc}| \leqslant |E(G)| + |E(G[T])| + |E_{cc}| < 2m$. This concludes the proof of the first assertion.

For the construction of $\mathcal{U}(G)$, we use the procedure described in Fig. 4 (some bookkeeping details are omitted). Step 1 adds vertex sets L and R (step 1.1) as well as those edges induced by them (step 1.2.1), and finds $N(T)$ (step 1.2.2). Step 2 adds edges in E_{cc} and E_c, and detect T_{cc} and T_c. Steps 3 and 4 add vertex w and edges incident to it. Step 5 cleans T. The dominating steps are 1 and 2, each of which checks every edge at most once, and hence the total time is $O(n + m)$. □

INPUT: a graph G and a hole H.
OUTPUT: the auxiliary graph $\mathcal{O}(G)$ or a forbidden induced subgraph of G.

0 **for each** $v \in V(G)$ **do** compute $\mathtt{first}(v)$ and $\mathtt{last}(v)$;
1 **for each** $v \in T$ **do**
1.1 add vertices v^l and v^r;
1.2 **for each** $u \in N(v)$ **do**
1.2.1 **if** $u \in T$ **then** add u^l to $N(v^l)$ and u^r to $N(v^r)$;
1.2.2 **else if** u is not marked **then** mark u and put it into $N(T)$;
2 **for each** $u \in N(T)$ **do**
2.1 **for each** $v \in N(u)$ **do** \\ $\mathtt{first}(v) \leqslant 0 \leqslant \mathtt{last}(v)$.
2.1.1 **if** $v \notin T$ **then goto** 2.1; \\ *henceforth* $v \in X$.
2.1.2 **if** $\mathtt{first}(v) = 0 = \mathtt{last}(v)$ **then return** a forbidden induced subgraph;
2.1.3 **if** $\mathtt{last}(v) = 0$ **then** replace v by v^l in $N(u)$ and add u to $N(v^l)$; \\ $uv \in E_{cc}$.
2.1.4 **if** $\mathtt{first}(v) = 0$ **then** replace v by v^r in $N(u)$ and add u to $N(v^r)$; \\ $uv \in E_c$.
\\ in the remaining cases $N[v] \cap V(H) = \{h_{-1}, h_0, h_1\}$ and $|N[u] \cap \{h_{-1}, h_1\}| = 1$.
2.1.5 **if** $\mathtt{last}(u) = |H| - 1$ **then** replace v by v^l in $N(u)$ and add u to $N(v^l)$; \\ $uv \in E_{cc}$.
2.1.6 **if** $\mathtt{first}(u) = 1$ **then** replace v by v^r in $N(u)$ and add u to $N(v^r)$; \\ $uv \in E_c$.
2.1.7 **if** $uv \in E_{cc}$ and v is not marked as T_{cc} **then** mark v and put it into T_{cc};
2.1.8 **if** $uv \in E_c$ and v is not marked as T_c **then** mark v and put it into T_c;
3 add vertex w;
4 **for each** $v \in T_{cc}$ **do** put w into $N(v^l)$ and v^l into $N(w)$;
5 remove T.

Figure 4: Procedure for constructing $\mathcal{O}(G)$ (Lemma 2.5).

In an interval model, each vertex v corresponds to a closed interval $I_v = [\mathtt{lp}(v), \mathtt{rp}(v)]$. Here $\mathtt{lp}(v)$ and $\mathtt{rp}(v)$ are the left and right *endpoints* of I_v respectively, and $\mathtt{lp}(v) < \mathtt{rp}(v)$. We use unit-length circles for circular-arc models, where every point has a positive value in $(0, 1]$. Each vertex v corresponds to a closed arc $A_v = [\mathtt{ccp}(v), \mathtt{cp}(v)]$. Here $\mathtt{ccp}(v)$ and $\mathtt{cp}(v)$ are counterclockwise and clockwise endpoints of A_v respectively; $0 < \mathtt{ccp}(v), \mathtt{cp}(v) \leqslant 1$ and they are assumed to be distinct. It is worth noting that possibly $\mathtt{cp}(v) < \mathtt{ccp}(v)$; such an arc necessarily contains the point 1.

The definition of $\mathcal{O}(G)$ is motivated by the following observation; see Fig. 5. We remark that this lemma is actually implied by Theorem 1.4, and the constructive proof presented below is to reveal the intuition.

Lemma 2.6. *If G is a normal Helly circular-arc graph, then $\mathcal{O}(G)$ is an interval graph.*

Proof. Let $\mathcal{A}$ be a normal Helly circular-arc model of G. Then the union of arcs $\{A_v : v \in T\}$ does not cover the circle. Assume that $\mathtt{ccp}(h_0) = 1$ and no other arc has an end at 1. An interval model of $\mathcal{O}(G)$ can be obtained from $\mathcal{A}$ by setting

- $I_{v^l} := [\mathtt{ccp}(v), \mathtt{cp}(v)]$ and $I_{v^r} := [\mathtt{ccp}(v) + 1, \mathtt{cp}(v) + 1]$ for $v \in T$ with $1 \notin A_v$;
- $I_{v^l} := [\mathtt{ccp}(v) - 1, \mathtt{cp}(v)]$ and $I_{v^r} := [\mathtt{ccp}(v), \mathtt{cp}(v) + 1]$ for $v \in T$ with $1 \in A_v$;
- $I_u := [\mathtt{ccp}(u), \mathtt{cp}(u)]$ for $u \in \overline{T}$; and
- $I_w := [-1, \max_{x \in \overline{T}} \mathtt{rp}(x) - 1]$.

It is easy to use the definition to verify that this gives an interval model of $\mathcal{O}(G)$. □

Note that for any vertex $v \in T$, an induced (v^l, v^r)-path corresponds to a cycle whose arcs cover the entire circle. The main thrust of our algorithm will be a process that does the reversed direction, which is nevertheless far more involved.

Theorem 2.7. *If $\mathcal{O}(G)$ is an interval graph, then we can in $O(n + m)$ time build a circular-arc model of G.*

Proof. We can in $O(n + m)$ time build an interval model $\mathcal{I}$ for $\mathcal{O}(G)$. By construction, $(w h_{-1}^l h_0^l h_1^l h_2 \cdots h_{-2} h_{-1}^r h_0^r h_1^r)$ is an induced path of $\mathcal{O}(G)$; without loss of generality, assume it goes "from left to right" in $\mathcal{I}$. We may assume $\mathtt{rp}(w) = 0$ and $\max_{u \in \overline{T}} \mathtt{rp}(u) = 1$, while no other interval in $\mathcal{I}$ has 0 or 1 as an endpoint. Let

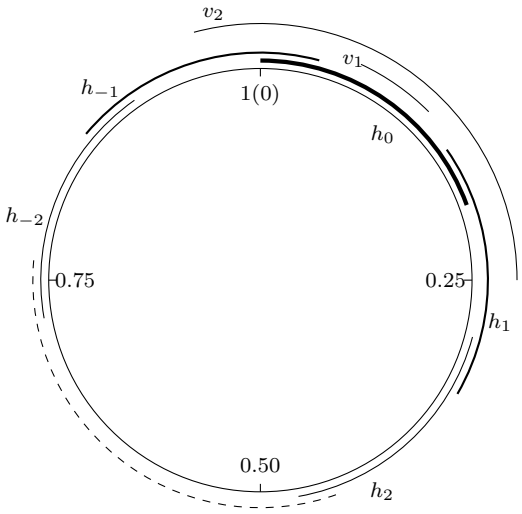

(a) A normal and Helly circular-arc model of a graph G.

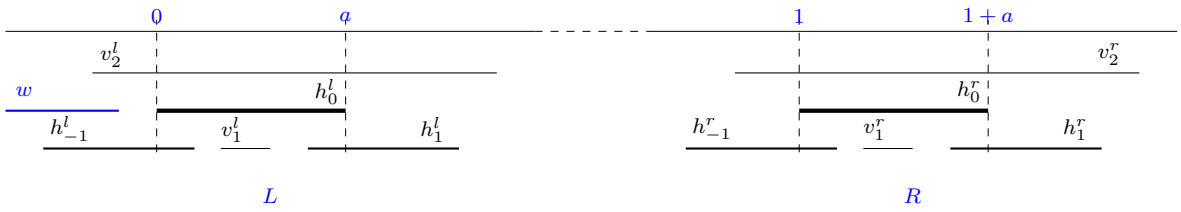

(b) An interval model of $\mathfrak{O}(G)$ derived from (a).

Figure 5: Illustration for Lemma 2.6.

$a = \mathtt{rp}(h_0^l)$. We use $\mathcal{I}$ to construct a set of arcs for $V(G)$ as follows. For each $u \in \overline{T}$, we set $A_u := [\mathtt{lp}(u), \mathtt{rp}(u)]$, which is a subset of the arc $(a, 1]$. For each $v \in T$, we set

$$A_v := \begin{cases} [\mathtt{lp}(v^r), \mathtt{rp}(v^l)] & \text{if } v \in T_{cc}, \\ [\mathtt{lp}(v^l), \mathtt{rp}(v^l)] & \text{otherwise.} \end{cases}$$

It remains to verify that the arcs obtained as such represent G, i.e., a pair of vertices u, v of G is adjacent if and only if A_u and A_v intersect. This holds trivially when neither u nor v is in T; hence we may assume without loss of generality that $v \in T$. By construction, $a < \mathtt{lp}(u) < \mathtt{rp}(u) \leqslant 1$ for every $u \in \overline{T}$. Note that $v^l \sim w$ and $v^r \sim \overline{T}$ for every $v \in T_{cc}$, which implies that $\mathtt{lp}(v^l) < 0$ if and only if $\mathtt{lp}(v^r) < 1$ if and only if $v \in T_{cc}$.

Assume first that u is also in T, then $u \sim v$ in G if and only if $u^l \sim v^l$ in $\mathfrak{O}(G)$. They are adjacent when both $u, v \in T_{cc}$, and since $\mathtt{lp}(v^l), \mathtt{lp}(u^l) < 0$, both A_u and A_v contains the point 1 and thus intersect. If neither u nor v is in T_{cc}, then $\mathtt{lp}(v^l), \mathtt{lp}(u^l) > 0$, and $u \sim v$ if and only if $A_u = [\mathtt{lp}(u^l), \mathtt{rp}(u^l)]$ and $A_v = [\mathtt{lp}(v^l), \mathtt{rp}(v^l)]$ intersect. Otherwise, assume, without loss of generality, that $\mathtt{lp}(v^l) < 0 < \mathtt{lp}(u^l)$, then $u \sim v$ in G if and only if $0 < \mathtt{lp}(u^l) < \mathtt{rp}(v^l)$, which implies A_u and A_v intersect (as both contain $[\mathtt{lp}(u^l), \mathtt{rp}(v^l)]$).

Assume now that u is not in T, and then $u \sim v$ in G if and only if either $u \sim v^l$ or $u \sim v^r$ in $\mathfrak{O}(G)$. In the case $u \sim v^l$, we have $\mathtt{lp}(v^l) \leqslant a < \mathtt{lp}(u) \leqslant \mathtt{rp}(v^l)$; since both A_u and A_v contain $[\mathtt{lp}(u), \mathtt{rp}(v^l)]$, which is nonempty, they intersect. In the case $u \sim v^r$, we have $\mathtt{lp}(v^r) < \mathtt{rp}(u) \leqslant 1$; since both A_u and A_v contain $[\mathtt{lp}(v^r), \mathtt{rp}(u)]$, they intersect. Otherwise, $u \not\sim v$ in G and $\mathtt{lp}(v^l) < \mathtt{rp}(v^l) < \mathtt{lp}(u) < \mathtt{rp}(u) < \mathtt{lp}(v^r) < \mathtt{rp}(v^r)$, then A_u and A_v are disjoint. □

We are now ready to present the recognition algorithm in Fig. 6, and prove Theorem 1.3. Recall that Lin et al. [10] have given a linear-time algorithm for verifying whether a circular-arc model is normal and Helly. In fact, given a circular-arc model, we can in linear time find the minimum number of arcs that cover the circle [2].

Algorithm **nhcag**(G)
Input: a graph G.
Output: either a normal Helly circular-arc model of G, or a forbidden induced subgraph of G.

1 test the chordality of G and find a hole H if not;
  if G is chordal **then** verify whether G is an interval graph or not;
2 construct the auxiliary grpah $\mathcal{U}(G)$;
3 **if** $\mathcal{U}(G)$ is not an interval graph **then**
    call Theorem 1.4 to find a forbidden induced subgraph;
4 **call** Theorem 2.7 to build a circular-arc $\mathcal{A}$ model of G;
5 verify whether $\mathcal{A}$ is normal and Helly.

Figure 6: The recognition algorithm for normal Helly circular-arc graphs.

Proof of Theorem 1.3. We use the algorithm nhcag presented in Fig. 6. Step 1 is clear. Steps 2 to 4 follow from Lemma 2.5, Theorem 1.4, and Lemma 2.7, respectively. If the model $\mathcal{A}$ built in step 4 is not normal and Helly, then we can in linear time find such a set of two or three arcs whose union covers the circle. Their corresponding vertices satisfy Lemma 2.4, and this concludes the proof. □

It is worth noting that if we are after a recognition algorithm (with positive certificate only), then we can simply return "NO" if the hypothesis of step 3 is true (see Lemma 2.6) and the algorithm is already complete.

3 Proof of Theorem 1.4

Recall that Theorem 1.4 is only called in step 3 of algorithm nhcag; the graph is then not choral and we have a hole H. In principle, we can pick any vertex as h_0. But for the convenience of presentation, we require it satisfies some additional conditions. If some vertex v is adjacent to four or more vertices in H, i.e., $\mathtt{last}(v)-\mathtt{first}(v) > 2$, then $v \notin H$. We can thus use $(h_{\mathtt{first}(v)} v h_{\mathtt{last}(v)})$ as a short cut for the sub-path induced by $N_H[v]$, thereby yielding a strictly shorter hole. This condition, that h_0 cannot be bypassed as such, is formally stated in the following lemma, and a procedure for finding such a hole is given in Fig. 7.

Lemma 3.1. *We can in $O(n+m)$ time find either a minimal forbidden induced subgraph, or a hole H such that $\{h_{-1}, h_0, h_1\} \subseteq N_H[v]$ for some v if and only if $N_H[v] = \{h_{-1}, h_0, h_1\}$.*

INPUT: a graph G and a hole H of G.
OUTPUT: a hole satisfying conditions of Lemma 3.1 or a minimal forbidden induced subgraph.

0 $h = h_0$; $a = -1$; $b = 1$;
1 **for each** $v \in V(G) \setminus H$ **do**
1.1 compute $\mathtt{first}(v)$ and $\mathtt{last}(v)$ in H;
1.2 **if** $(\mathtt{first}(v) < a$ and $\mathtt{last}(v) \geqslant b)$ or $(\mathtt{first}(v) = a$ and $\mathtt{last}(v) > b)$ **then**
       $h = v$; $a = \mathtt{first}(v)$; $b = \mathtt{last}(v)$; \\ $\mathtt{first}(v) < 0 < \mathtt{last}(v)$.
2 **return** $(h h_b h_{b+1} \cdots h_a h)$ where h is the new h_0.

Figure 7: Procedure for finding the hole for Lemma 3.1.

Proof. We apply the procedure given in Fig. 7. Step 1 greedily searches for an inclusion-wise maximal $N_H[v]$ satisfying $\mathtt{first}(v) \leqslant -1$ and $1 \leqslant \mathtt{last}(v)$. Initially, $a = \mathtt{first}(h_0) = -1$ and $b = \mathtt{last}(h_0) = 1$. Each iteration of step 1 checks an unexplored vertex v in $V(G) \setminus H$. If either condition of step 1.2 is satisfied, then $N[v]$ properly contains $\{h_a, h_{a+1}, \ldots, h_b\}$, and a and b are updated to be $\mathtt{first}(v)$ and $\mathtt{last}(v)$ respectively. Note that the values of a and b are non-increasing and nondecreasing respectively. Thus, no previously explored vertex is adjacent to all of $\{h_a, h_{a+1}, \ldots, h_b\}$. After step 1, all vertices have been explored, and the hole $(h h_b h_{b+1} \cdots h_a h)$ satisfies the claimed condition.

What dominates the procedure is finding $\texttt{first}(v)$ and $\texttt{last}(v)$ for all vertices (step 1.1). It takes $O(d(v))$ time for each vertex v and $O(n+m)$ time in total. □

This linear-time procedure can be called before step 2 of algorithm nhcag, and it does not impact the asymptotic time complexity of the algorithm, which remains linear. Henceforth we may assume that H satisfies the condition of Lemma 3.1. During the construction of $\mathfrak{V}(G)$, we have checked $N_H[v]$ for every vertex v, and Lemma 2.2 was called if it applies. Therefore, for the proof of Theorem 1.4 in this section, we may assume that $N_H[v]$ always induces a proper sub-path of H.

Each vertex x of $\mathfrak{V}(G)$ different from w is uniquely defined by a vertex of G, which is denoted by $\phi(x)$. We say that x is *derived from* $\phi(x)$. For example, $\phi(v^l) = \phi(v^r) = v$ for $v \in T$. By abuse of notation, we will use the same letter for a vertex $u \in \overline{T}$ of G and the unique vertex of $\mathfrak{V}(G)$ derived from u; its meaning is always clear from the context. Therefore, $\phi(u) = u$ for $u \in \overline{T}$, and in particular, $\phi(h_i) = h_i$ for $i = 2, \ldots, |H|-2$. We can mark $\phi(x)$ for each vertex of $\mathfrak{V}(G)$ during its construction. The function ϕ is also generalized to a set U of vertices that does not contain w, i.e., $\phi(U) = \{\phi(v) : v \in U\}$. We point out that possibly $|\phi(U)| \neq |U|$.

By construction of $\mathfrak{V}(G)$, if a pair of vertices x and y (different from w) is adjacent in $\mathfrak{V}(G)$, then $\phi(x)$ and $\phi(y)$ must be adjacent in G as well. The converse is not necessarily true, e.g., for any vertex $v \in T_c$ and edge $uv \in E_c$, we have $u \not\sim v^r$, and for any pair of adjacent vertices $u, v \in T$, we have $u^l \not\sim v^r$ and $u^r \not\sim v^l$. We say that a pair of vertices x, y of $\mathfrak{V}(G)$ is a *bad pair* if $\phi(x) \sim \phi(y)$ in G but $x \not\sim y$ in $\mathfrak{V}(G)$. By definition, w does not participate in any bad pair, and at least one vertex of a bad pair is in $L \cup R$. Note that any induced path of length d between a bad pair x, y with $x = v^l$ or v^r can be extended to a (v^l, v^r)-path with length $d+1$.

We have seen that if G is a normal Helly circular-arc graph, then for any $v \in T$, the distance between v^l and v^r is at least 4. We now see what happens when this necessary condition is not satisfied by $\mathfrak{V}(G)$. By definition of $\mathfrak{V}(G)$, there is no edge between L and R; for any $v \in T$, there is no vertex adjacent to both v^l and v^r. In other words, for every $v \in T$, the distance between v^l and v^r is at least 3. The following observation can be derived from Lemmas. 2.2 and 2.3.

Lemma 3.2. *Given a (v^l, v^r)-path P of length 3 for some $v \in T$, we can in $O(n+m)$ time find a minimal forbidden induced subgraph.*

Proof. Let $P = (v^l x y v^r)$. Note that P must be a shortest (v^l, v^r)-path, and $w \notin P$. The inner vertices x and y cannot be both in $L \cup R$; without loss of generality, let $x \in \overline{T}$. Assume first that $y \in \overline{T}$ as well, i.e., $vx \in E_c$ and $vy \in E_{cc}$. By definition, $v \in T_c \cap T_{cc}$, and then v is adjacent to both h_{-1} and h_1. If follows from Lemma 3.1 that $N_H[v] = \{h_{-1}, h_0, h_1\}$, and then $x \sim h_1$ and $y \sim h_{-1}$. If $x \sim h_{-1}$, i.e., $\texttt{last}(x) = |H|-1$, then we call Lemma 2.3 with v and x. If $\texttt{last}(x) < \texttt{first}(y)$, then we call Lemma 2.2 with x and y. In the remaining case, $\texttt{first}(y) \leqslant \texttt{last}(x) < |H|-1$, and $(vxh_{\texttt{last}(x)} \cdots h_{-1}v)$ is a hole of G; this hole is completely adjacent to y, and thus we find a wheel.

Now assume that, without loss of generality, $y = u^r \in R$. If $\texttt{last}(v) \geqslant \texttt{first}(y)$, then we call Lemma 2.3 with v and y. Otherwise, $(vh_{\texttt{last}(v)} \cdots h_{\texttt{first}(y)}uv)$ is a hole of G; this hole is completely adjacent to x, and thus we find a wheel. □

If G is a normal Helly circular-arc graph, then in a circular-arc model of G, all arcs for T_{cc} and T_c contain $\texttt{ccp}(h_0)$ and $\texttt{cp}(h_0)$ respectively. Thus, both T_{cc} and T_c induce cliques. This observation is complemented by the following lemma.

Lemma 3.3. *Given a pair of nonadjacent vertices $u, x \in T_{cc}$ (or T_c), we can in $O(n+m)$ time find a minimal forbidden induced subgraph of G.*

Proof. By definition, we can find edges $uv, xy \in E_{cc}$. We have three (possibly intersecting) induced paths $h_0h_1h_2$, h_0uv, and h_0xy. If both u and x are adjacent to h_1, then we return $(uh_{-1}xh_1u)+h_0$ as a wheel. Hence we may assume $x \not\sim h_1$.

If $u \sim h_1$, then by Lemma 3.1, $N_H[u] = \{h_{-1}, h_0, h_1\}$. We consider the subgraph induced by the set of distinct vertices $\{h_0, h_1, h_2, u, v, x\}$. If v is adjacent to h_0 or h_1, then we can call Lemma 2.4 with u and v. By assumption, h_0, h_1, and u make a triangle; x is adjacent to neither u nor h_1; and h_2 is adjacent to neither h_0 nor u. Therefore, the only uncertain adjacencies in this subgraph are between v, x, and h_2. The subgraph is thus isomorphic to (1) FIS-1 if there are two edges among v, x, and h_2; (2) $\overline{C_6}$ if v, x, and h_2 are pairwise adjacent; or (3) net if v, x, and h_2 are pairwise nonadjacent. In the remaining cases there is precisely one edge among v, x, and h_2, then we can return a C^*, e.g., (vxh_0uv) and h_2 when the edge is vx.

Assume now that u, x, and h_1 are pairwise nonadjacent. We consider the subgraph induced by $\{h_0, h_1, h_2, u, v, x, y\}$, where the only uncertain relations are between v, y, and h_2. The subgraph is thus isomorphic to (1) $K_{2,3}$ if all of them are identical; or (2) twin-C_5 if two of them are identical, and adjacent to the other. If two of them are identical, and nonadjacent to the other, then the subgraph contains a C^*, e.g., (vuh_0xv) and h_2 when $v = y$. In the remaining cases, all of v, y, and h_2 are distinct, and then the subgraph (1) is isomorphic to long claw if they are pairwise nonadjacent; (2) contains net $\{h_1, h_2, u, v, x, y\}$ if they are pairwise adjacent; or (3) is isomorphic to FIS-2 if there are two edges among them. If there is one edge among them, then the subgraph contains a C^*, e.g., (vuh_0xyv) and h_2 when the edge is vy.

A symmetrical argument applies to T_c. Edges uv and xy can be found in $O(n)$ time, and only a small constant number of adjacencies are checked; it thus takes $O(n + m)$ time in total. □

It can be checked in linear time whether T_{cc} and T_c induce cliques. When it is not, a pair of nonadjacent vertices can be found in the same time. By Lemma 3.3, we may assume hereafter that T_{cc} and T_c induce cliques. We say that a vertex v is *simplicial* if $N[v]$ induces a clique. Recall that $N(w) \subseteq T_{cc}$; as a result, w is simplicial and participates in no holes.

Proposition 3.4. *Given an (h_0^l, h_0^r)-path P that is nonadjacent to h_i for some $1 < i < |H| - 1$, we can in $O(n + m)$ time find a minimal forbidden induced subgraph.*

Proof. Inside P there must be a sub-path P' whose ends x, y are in L and R respectively, and whose inner vertices are all in $\overline{T}$. Let x' and y' be the neighbors of x and y in P' respectively; note that they are both in $\overline{T}$. If $N_H[\phi(x)]$ and $N_H[x']$ are disjoint, then we call Lemma 2.3. Otherwise, by assumption, we have $0 < \texttt{first}(x') \leqslant \texttt{last}(x') < i$; likewise, we may assume $i < \texttt{first}(y') \leqslant \texttt{last}(y') < |H|$. Starting from x', we traverse P' till the first pair of consecutive vertices u, v in P such that $N_H[u]$ and $N_H[v]$ are disjoint: note that such a pair must exist because no vertex between x' and y' is adjacent to h_0 or h_i. Then we call Lemma 2.3. □

We are now ready to prove Theorem 1.4, which is separated into three statements, the first of which considers the case when $\mathcal{U}(G)$ is not chordal.

Lemma 3.5. *Given a hole C of $\mathcal{U}(G)$, we can in $O(n + m)$ time find a minimal forbidden induced subgraph.*

Proof. Let us first take care of some trivial cases. If C is contained in L or R or $\overline{T}$, then by construction, $\phi(C)$ is a hole of G. This hole is either nonadjacent or completely adjacent to h_0 in G, whereupon we can return $\phi(C)$ and h_0 as a C^* or wheel respectively. Since L and R are nonadjacent, one of the cases above must hold if C is disjoint from $\overline{T}$. Henceforth we may assume that C intersects $\overline{T}$ and, without loss of generality, L; it might intersect R as well, but this fact is irrelevant in the following argument. Then we can find an edge x_1x_2 of C such that $x_1 \in L$ and $x_2 \in \overline{T}$, i.e., $x_1x_2 \in E_c$.

Let $a := \texttt{last}(\phi(x_1))$. Assume first that $x_2 = h_a$; then we must have $a > 1$. Let x_3 and x_4 be the next two vertices of C. Note that $x_3 \notin L$, i.e., $x_3 \not\sim h_0^l$; otherwise $x_1 \sim x_3$, which is impossible. If $x_3 \sim h_{a-2}$ (or h_{a-2}^l when $a = 3$), then $\phi(\{x_1, x_2, x_3\}) \cup \{h_{a-2}\}$ induces a hole of G, and we can return it and h_{a-1} as a wheel. Note that $x_4 \not\sim h_a$ as they are non-consecutive vertices of the hole C. We now argue that $\texttt{last}(\phi(x_4)) < a$. Suppose for contradiction, $\texttt{first}(\phi(x_4)) > a$. We can extend the (x_3, x_1)-path P in C that avoids x_2 to a (h_0^l, h_0^r)-path avoiding the neighborhood of h_a, which allows us to call Proposition 3.4. We can call Lemma 2.3 with x_3 and x_4 if $\texttt{first}(\phi(x_3)) = a$. In the remaining case, $\texttt{first}(\phi(x_3)) = a - 1$. Let x be the first vertex in P that is adjacent to h_{a-2} (or h_{a-2}^l if $a \leqslant 3$); its existence is clear as x_1 satisfies this condition. Then $\phi(\{x_3, \ldots, x, h_{a-2}, x_1, x_2\})$ induces a hole of G, and we can return it and h_{a-1} as a wheel.

Assume now that h_a is not in C. Denote by P the (x_2, x_1)-path obtained from C by deleting the edge x_1x_2. Let x be the first neighbor of h_{a+1} in P, and let y be either the first neighbor of h_{a-1} in the (x, x_1)-path or the other neighbor of x_1 in C. It is easy to verify that $\phi(\{x_1, \cdots, x, \cdots y, x_2\})$ induces a hole of G, which is completely adjacent to h_a, i.e., we have a wheel. □

In the rest $\mathcal{U}(G)$ will be chordal, and thus we have a chordal minimally non-interval subgraph F of $\mathcal{U}(G)$. This subgraph is isomorphic to some graph in Fig. 1, on which we use the following notation. It is immediate from Fig. 1 that each of them contains precisely three simplicial vertices (squared vertices), which are called *terminals*, and others (round vertices) are *non-terminal vertices*. In a long claw or †, for each $i = 1, 2, 3$, terminal t_i has a unique neighbor, denoted by u_i.

Since the diameter and maximum clique of F is at most four, all bad pairs in it can be found easily.

Proposition 3.6. *Given a subgraph* F *of* $U(G)$ *in Fig. 1, we can in* $O(n+m)$ *time find either all bad pairs in* F *or a minimal forbidden induced subgraph.*

Proof. A long claw or whipping top has only seven vertices, and thus will not concern us. A bad pair is either between L and R, or between $\{v^l : v \in T_{cc}\} \cup \{v^r : v \in T_c\}$ and $\overline{T}$. Let $L_c = \{v^l : v \in T_c\}$ and $R_{cc} = \{v^r : v \in T_{cc}\}$; both induce cliques. Since a clique of F contains at most 4 vertices, it contains at most 8 vertices of L_c and R_{cc}. Bad pairs intersecting them can thus be found in linear time. We now consider other bad pairs, which must be between $L \setminus L_c$ and $R \setminus R_{cc}$. By construction, there is no edge between $L \setminus L_c$ and $R \cup \overline{T}$; there is no edge between $R \setminus R_{cc}$ and $L \cup \overline{T}$. Therefore, the shortest distance between $L \setminus L_c$ and $R \setminus R_{cc}$ is at least 4. There exists only one pair of distance 4 vertices in a †, and no such a pair in a ‡. □

Lemma 3.7. *Given a subgraph* F *of* $U(G)$ *in Fig. 1 that does not contain* w*, we can in* $O(n+m)$ *time find a minimal forbidden induced subgraph.*

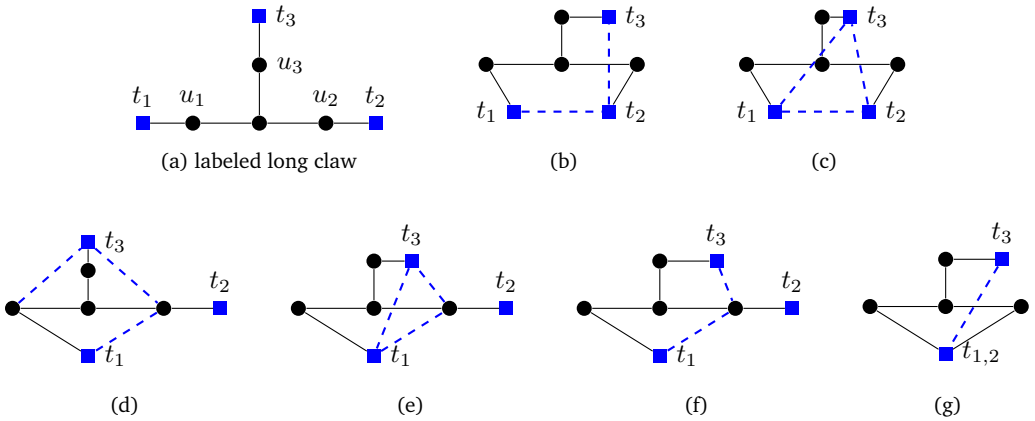

Figure 8: Illustrations for Lemma 3.7 (blue dashed edges are in G only).

Proof. We first call Proposition 3.6 to find all bad pairs in F. If F has no bad pair, then we return the subgraph of G induced by $\phi(F)$, which is isomorphic to F. Let x, y be a bad pair with the minimum distance in F; we may assume that it is 3 or 4, as otherwise we can call Lemma 3.2. Noting that the distance between a pair of non-terminal vertices is at most 2, we may assume that without loss of generality, x is a terminal of F. We break the argument based on the type of F.

Long claw. We may assume that $x = t_1$ and $y \in \{u_2, t_2\}$; other situations are symmetrical. Let P be the unique (x, y)-path in F. If $\phi(t_3)$ is nonadjacent to $\phi(P)$, then we return $\phi(P)$ and $\phi(t_3)$ as a C^*; we are thus focused on the adjacency between $\phi(t_3)$ and $\phi(P)$ (Fig. 8). If $y = t_2$, then by the selection of x, y (they have the minimum distance among all bad pairs), $\phi(t_3)$ can be only adjacent to $\phi(t_1)$ and/or $\phi(t_2)$. We return either $\phi(F)$ as an FIS-2 (Fig. 8b), or $\phi(\{t_1, t_2, t_3, u_1, u_2, u_3\})$ as a net (Fig. 8c). In the remaining cases, $y = u_2$, and $\phi(t_3)$ can only be adjacent to $\phi(u_1)$, $\phi(u_2)$, and/or $\phi(t_1)$. We point out that possibly $\phi(t_2) = \phi(t_1)$, which is irrelevant as $\phi(t_2)$ will not be used below. If $\phi(t_3)$ is adjacent to both $\phi(u_1)$ and $\phi(u_2)$ in G, then we get a $K_{2,3}$ (Fig. 8d). Note that this is the only case when $\phi(t_1) = \phi(t_3)$. If $\phi(t_3)$ is adjacent to both $\phi(t_1)$ and $\phi(u_2)$ in G, then we get an FIS-1 (Fig. 8e). If $\phi(t_3)$ is adjacent to only $\phi(u_2)$ or only $\phi(t_1)$ in G, then we get a domino (Fig. 8f) or twin-C_5 (Fig. 8g), respectively. The situation that $\phi(t_3)$ is adjacent to $\phi(u_1)$ but not $\phi(u_2)$ is similar as above.

Whipping top. The diameter is 3, and this distance is attained only by $\{t_1, t_3\}$ or $\{t_2, t_3\}$. If both are bad pairs, then we have a domino. If $\{t_1, t_3\}$ is the only bad pair, then $\phi(F \setminus N[t_2])$ induces a hole of G, and it is nonadjacent to $\phi(t_2)$; we get a C^*. A symmetrical argument applies if $\{t_2, t_3\}$ is the only bad pair.

†. Consider first that $x = t_1$ and $y = t_3$, and let $P = (t_1 u_1 u_3 t_3)$. If $\phi(t_2)$ is nonadjacent to the hole induced by $\phi(P)$, then we return $\phi(P)$ and $\phi(t_2)$ as a C^*. If $\phi(t_2)$ is adjacent to $\phi(t_3)$ or $\phi(u_1)$, then we get a domino. If $\phi(t_2)$ is adjacent to $\phi(t_1)$, then we get a twin-C_5. If $\phi(t_2)$ is adjacent to $\phi(t_1)$ and precisely one of $\{\phi(t_3), \phi(u_1)\}$, then we get an FIS-1. If $\phi(t_2)$ is adjacent to both $\phi(t_3)$ and $\phi(u_1)$, then we get a $K_{2,3}$; here the adjacency between $\phi(t_2)$ and $\phi(t_1)$ is immaterial. A symmetric argument applies when $\{t_2, t_3\}$ is a bad pair. In the remaining

case, neither $\phi(t_1)$ nor $\phi(t_2)$ is adjacent to $\phi(t_3)$. Therefore, a bad pair must be in the path $F - N[t_3]$, which is nonadjacent to $\phi(t_3)$, then we get a C^*.

‡. The only pair of vertices of distance 3 is $\{t_1, t_2\}$. Let P be the (t_1, t_2)-path in $F - N[t_3]$. Since $\phi(t_3)$ cannot be adjacent to any vertex in $\phi(P)$, we can return $\phi(P)$ and $\phi(t_3)$ as a C^*. □

Lemma 3.8. *Given a subgraph F of $\mathcal{U}(G)$ in Fig. 1 that contains w, we can in $O(n+m)$ time find a minimal forbidden induced subgraph.*

\\ Note that $0 \leqslant \texttt{last}(\phi(x_1)), \texttt{last}(\phi(x_2)) \leqslant 1$.
1 **if** $\texttt{last}(\phi(x_1)) = 1$ and $y_1 \sim h_2$ **then**
call Lemma 2.3 with $(y_1\phi(x_1)h_1h_2y_1)$ and $\{\phi(x_2), y_2\}$;
if $\texttt{last}(\phi(x_1)) = 0$ and $y_1 \sim h_1$ **then**
call Lemma 2.3 with $(y_1\phi(x_1)h_0h_1y_1)$ and $\{\phi(x_2), y_2\}$;
if $y_2 \sim h_{\texttt{last}(\phi(x_2))+1}$ **then** symmetric as above;
2 **if** $\texttt{last}(\phi(x_1)) = \texttt{last}(\phi(x_2))$ **then**
return $\{y_1, \phi(x_1), y_2, \phi(x_2), h_{\texttt{last}(\phi(x_2))}, h_{\texttt{last}(\phi(x_2))+1}\}$ as a †; \\ *Fig. 10a.*
\\ assume from now that $\texttt{last}(\phi(x_1)) = 1$ and $\texttt{last}(\phi(x_2)) = 0$.
3 **if** $\phi(x_2) \sim h_2$ **then return** $(\phi(x_2)h_0h_1h_2\phi(x_2))$ and y_1 as a C^*;
4 **if** $y_2 \not\sim h_{-1}$ **then return** $\{y_1, h_{-1}, \phi(x_1), y_2, \phi(x_2), h_0, h_1\}$ as a ‡; \\ *Fig. 10b.*
if $y_2 \sim h_{-1}$ **then return** $\{y_1, h_{-1}, y_2, \phi(x_2), h_0, h_1\}$ as a †. \\ *Fig. 10c.*

Figure 9: Procedure for Lemma 3.8

Proof. Since w is simplicial, it has at most 2 neighbors in F. If w has a unique neighbor in F, then we can use a similar argument as Lemma 3.7. Now let x_1, x_2 be the two neighbors of w in F. If there exists some vertex $u \in \overline{T}$ adjacent to both $\phi(x_1)$ and $\phi(x_2)$ in G, which can be found in linear time, then we can use it to replace w. Hence we assume there exists no such vertex. By assumption, we can find two distinct vertices $y_1, y_2 \in \overline{T}$ such that $\phi(x_1)y_1, \phi(x_2)y_2 \in E_{cc}$; note that $\phi(x_1) \not\sim y_2$ and $\phi(x_2) \not\sim y_1$ in G. As a result, y_1 and y_2 are nonadjacent; otherwise, $\{y_1, y_2\}$ and the counterparts of $\{x_1, x_2\}$ in R induce a hole of $\mathcal{U}(G)$, which contradicts the assumption that $\mathcal{U}(G)$ is chordal. We then apply the procedure described in Fig. 9.

We now verify the correctness of the procedure. Since each step—either directly or by calling a previously verified lemma—returns a minimal forbidden induced subgraph of G, all conditions of previous steps are assumed to not hold in a later step. By Lemma 3.1, $\texttt{last}(\phi(x_1))$ and $\texttt{last}(\phi(x_2))$ are either 0 or 1. Step 1 considers the case where $y_1 \sim h_{\texttt{last}(\phi(x_1))+1}$. By Lemma 2.4, $y_1 \not\sim h_{\texttt{last}(\phi(x_1))}$. Thus, $(y_1\phi(x_1)h_1h_2y_1)$ or $(y_1\phi(x_1)h_0h_1y_1)$ is a hole of G, depending on $\texttt{last}(\phi(x_1))$ is 0 or 1. In the case $(y_1\phi(x_1)h_1h_2y_1)$, only $\phi(x_1)$ and h_1 can be adjacent to $\phi(x_2)$; they are nonadjacent to y_2. Likewise, in the case $(y_1\phi(x_1)h_0h_1y_1)$, vertices $\phi(x_1)$ and h_0 are adjacent to $\phi(x_2)$ but not y_2, while h_1 can be adjacent to only one of $\phi(x_2)$ and y_2. Thus, we can call Lemma 2.3. A symmetric argument applies when $y_2 \sim h_{\texttt{last}(\phi(x_2))+1}$. Now that the conditions of step 1 do not hold true, step 2 is clear from assumption. Henceforth we may assume without loss of generality that $\texttt{last}(\phi(x_1)) = 1$ and $\texttt{last}(\phi(x_2)) = 0$. Consequently, $\texttt{last}(y_1) = |H| - 1$ (Lemma 2.3). Because we assume that the condition of step 1 does not hold, $y_1 \not\sim h_2$; this justifies step 3. Step 4 is clear as y_1 is always adjacent to h_{-1}. □

References

[1] Yixin Cao. Direct and certifying recognition of normal Helly circular-arc graphs in linear time. In Jianer Chen, John Hopcroft, and Jianxin Wang, editors, *Proceedings of the 8th International Frontiers of Algorithmics Workshop, FAW 2014*, 2014. To appear.

[2] Yixin Cao. Linear recognition of almost (unit) interval graphs. arXiv:1403.1515, 2014.

[3] Guillermo Durán, Luciano N. Grippo, and Martín D. Safe. Structural results on circular-arc graphs and circle graphs: A survey and the main open problems. *Discrete Applied Mathematics*, 164(2):427–443, 2014. LAGOS'11: Sixth Latin American Algorithms, Graphs, and Optimization Symposium, Bariloche, Argentina—2011.

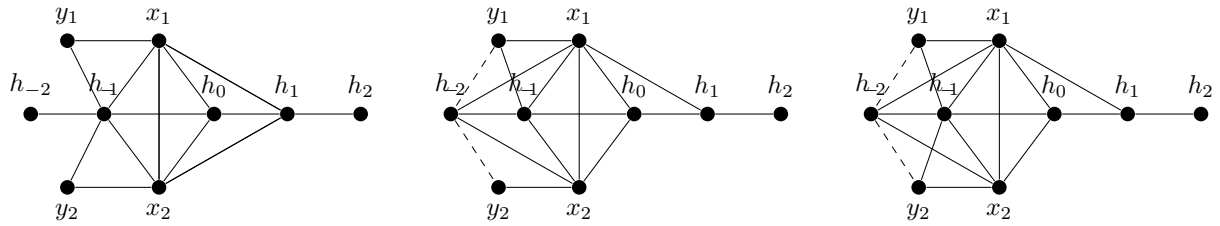

Figure 10: Structures used in the proof of Lemma 3.8 (dashed edges may or may not exist).

[4] Luciano N. Grippo and Martín D. Safe. On circular-arc graphs having a model with no three arcs covering the circle. CLAIO-SBPO 2012, Rio de Janeiro - Brazil, September, 24-28 2012. http://www2.claiosbpo2012.iltc.br/pdf/102012.pdf.

[5] Hugo Hadwiger, Hans Debrunner, and Victor Klee. *Combinatorial geometry in the plane*. Athena series. Holt, Rinehart and Winston, London, 1964.

[6] Pinar Heggernes and Dieter Kratsch. Linear-time certifying recognition algorithms and forbidden induced subgraphs. *Nordic Journal of Computing*, 14(1-2):87–108, 2007.

[7] Dieter Kratsch, Ross M. McConnell, Kurt Mehlhorn, and Jeremy P. Spinrad. Certifying algorithms for recognizing interval graphs and permutation graphs. *SIAM Journal on Computing*, 36(2):326–353, 2006. A preliminary version appeared in SODA 2003.

[8] Cornelis G. Lekkerkerker and J. Ch. Boland. Representation of a finite graph by a set of intervals on the real line. *Fundamenta Mathematicae*, 51:45–64, 1962.

[9] Min Chih Lin, Francisco J. Soulignac, and Jayme L. Szwarcfiter. The clique operator on circular-arc graphs. *Discrete Applied Mathematics*, 158(12):1259–1267, 2010.

[10] Min Chih Lin, Francisco J. Soulignac, and Jayme L. Szwarcfiter. Normal Helly circular-arc graphs and its subclasses. *Discrete Applied Mathematics*, 161(7-8):1037–1059, 2013.

[11] Min Chih Lin and Jayme L. Szwarcfiter. Characterizations and recognition of circular-arc graphs and subclasses: A survey. *Discrete Mathematics*, 309(18):5618–5635, 2009.

[12] Nathan Lindzey and Ross M. McConnell. On finding Tucker submatrices and Lekkerkerker-Boland subgraphs. In Andreas Brandstädt, Klaus Jansen, and Rüdiger Reischuk, editors, *Revised Papers of the 39th International Workshop on Graph-Theoretic Concepts in Computer Science, WG 2013*, volume 8165 of *LNCS*, pages 345–357, 2013.

[13] Ross M. McConnell. Linear-time recognition of circular-arc graphs. *Algorithmica*, 37(2):93–147, 2003.

[14] Ross M. McConnell, Kurt Mehlhorn, Stefan Näher, and Pascal Schweitzer. Certifying algorithms. *Computer Science Review*, 5(2):119–161, 2011.

[15] Terry A. McKee. Restricted circular-arc graphs and clique cycles. *Discrete Mathematics*, 263(1-3):221–231, 2003.

[16] Robert E. Tarjan and Mihalis Yannakakis. Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. *SIAM Journal on Computing*, 13(3):566–579, 1984. With Addendum in the same journal, 14(1):254-255, 1985.

[17] Alan C. Tucker. Structure theorems for some circular-arc graphs. *Discrete Mathematics*, 7(1-2):167–195, 1974.

[18] Alan C. Tucker. Coloring a family of circular arcs. *SIAM Journal on Applied Mathematics*, 29(3):493–502, 1975.